New Hires – “Wisdom of the Ages(d)”

Make new friends, but keep the old;
Those are silver, these are gold.

New-made friendships, like new wine,
Age will mellow and refine.

Friendships that have stood the test -
Time and change - are surely best;

Brow may wrinkle, hair grow gray,
Friendship never knows decay.

Make new friends, but keep the old;
Those are silver, these are gold.

“It may be that your sole purpose in life is to serve as a warning to others”

Modeling Bird Migration under Climate Change – A Mechanistic Approach

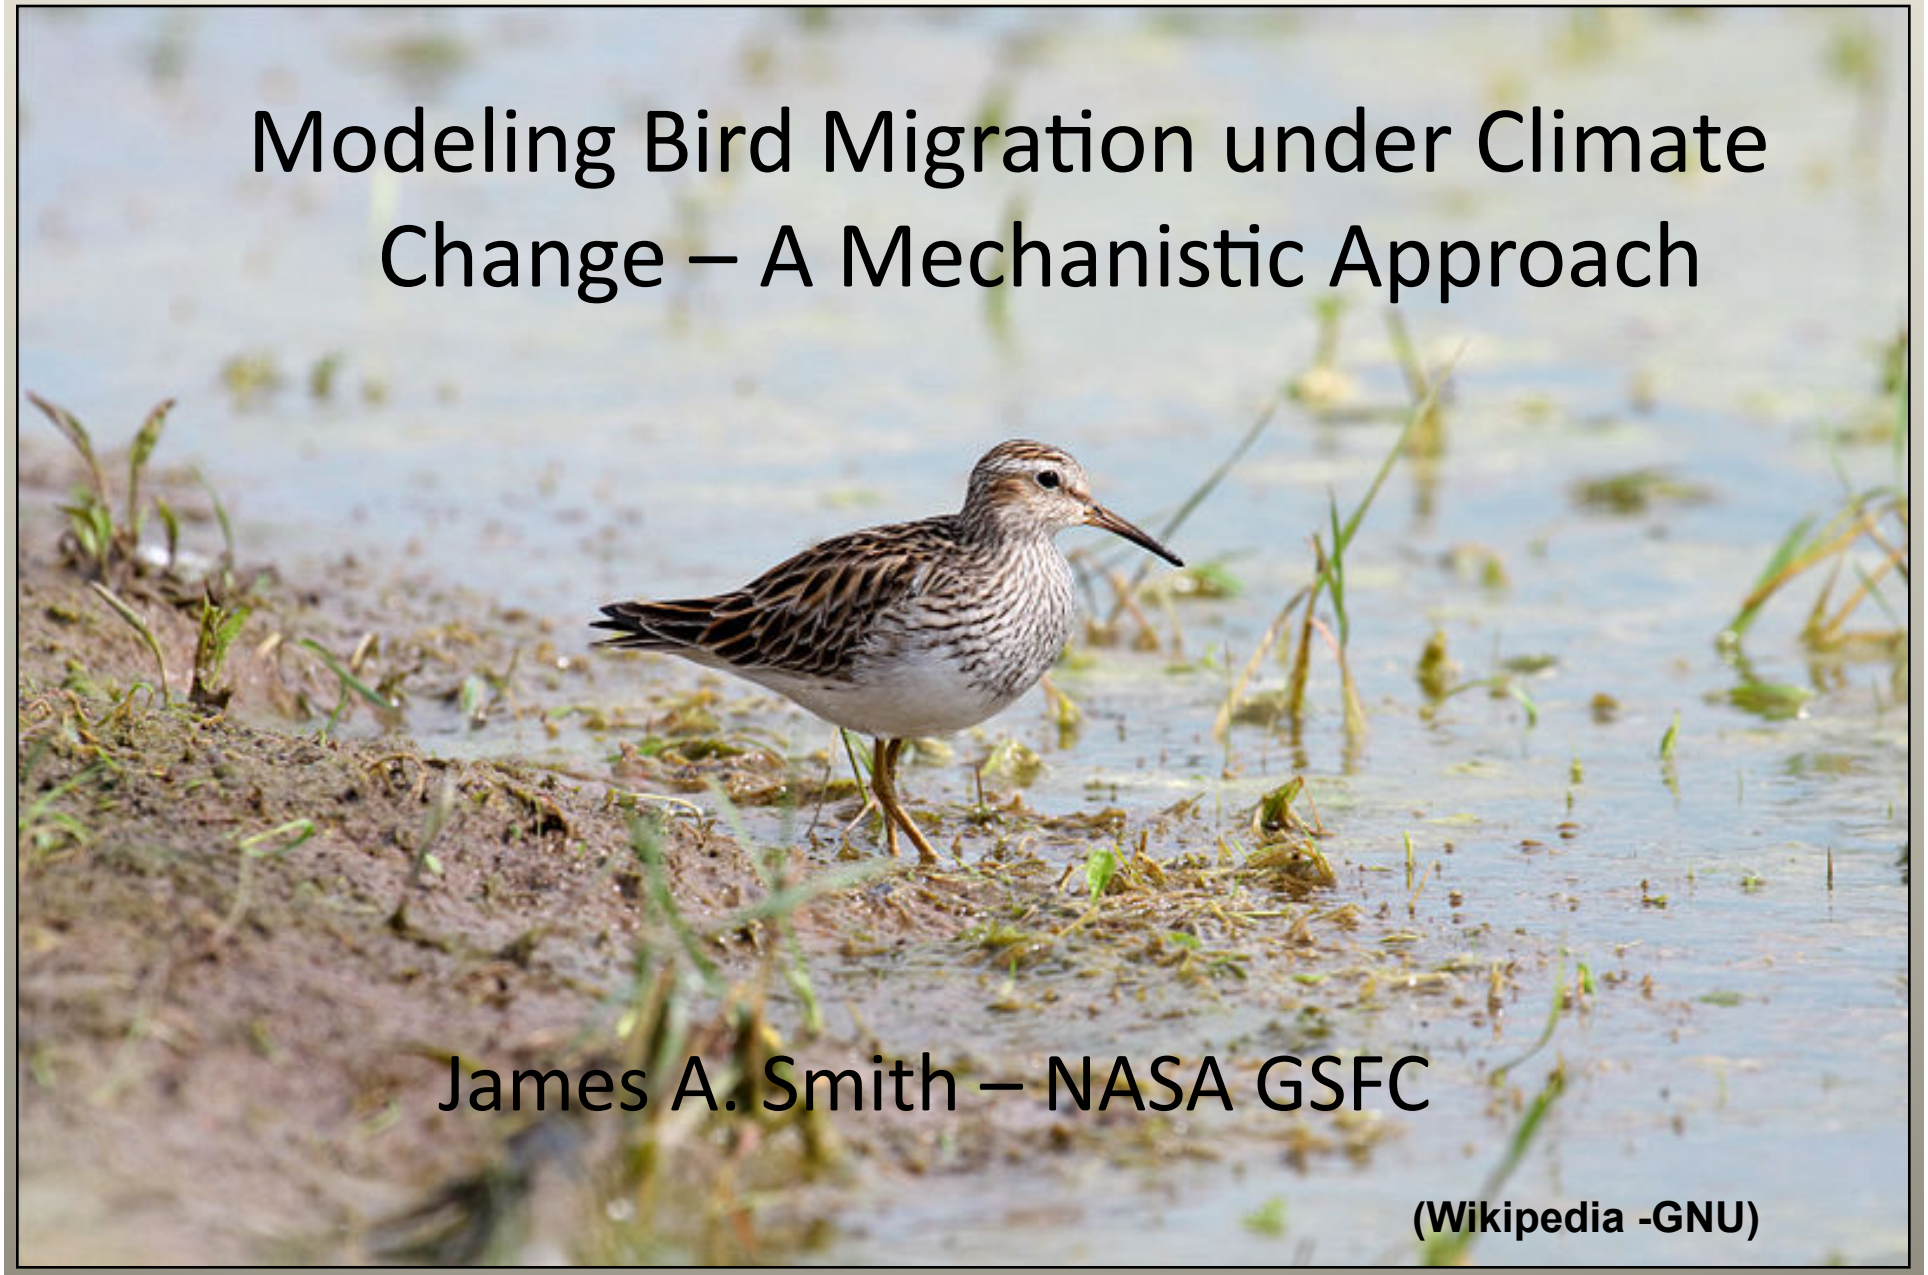

James A. Smith – NASA GSFC

(Wikipedia -GNU)

Goal

Use physics and biologically based models to Understand how migrating organisms respond to changes in their environment –

What are the impacts of resulting changes in the quality, location, and quantity of stopover habitat?

What is the coupling between timing of migration and key environmental processes?

With Environmental Change

Will they still have a leg to stand on?

Model

- Biophysical flight model (avian energetics)
- Movement behavior and decision rules
- Daily time step
- Arbitrary geographic grid

Simulate the migration routes, timing and energy budgets of individual birds under dynamic weather and land surface conditions

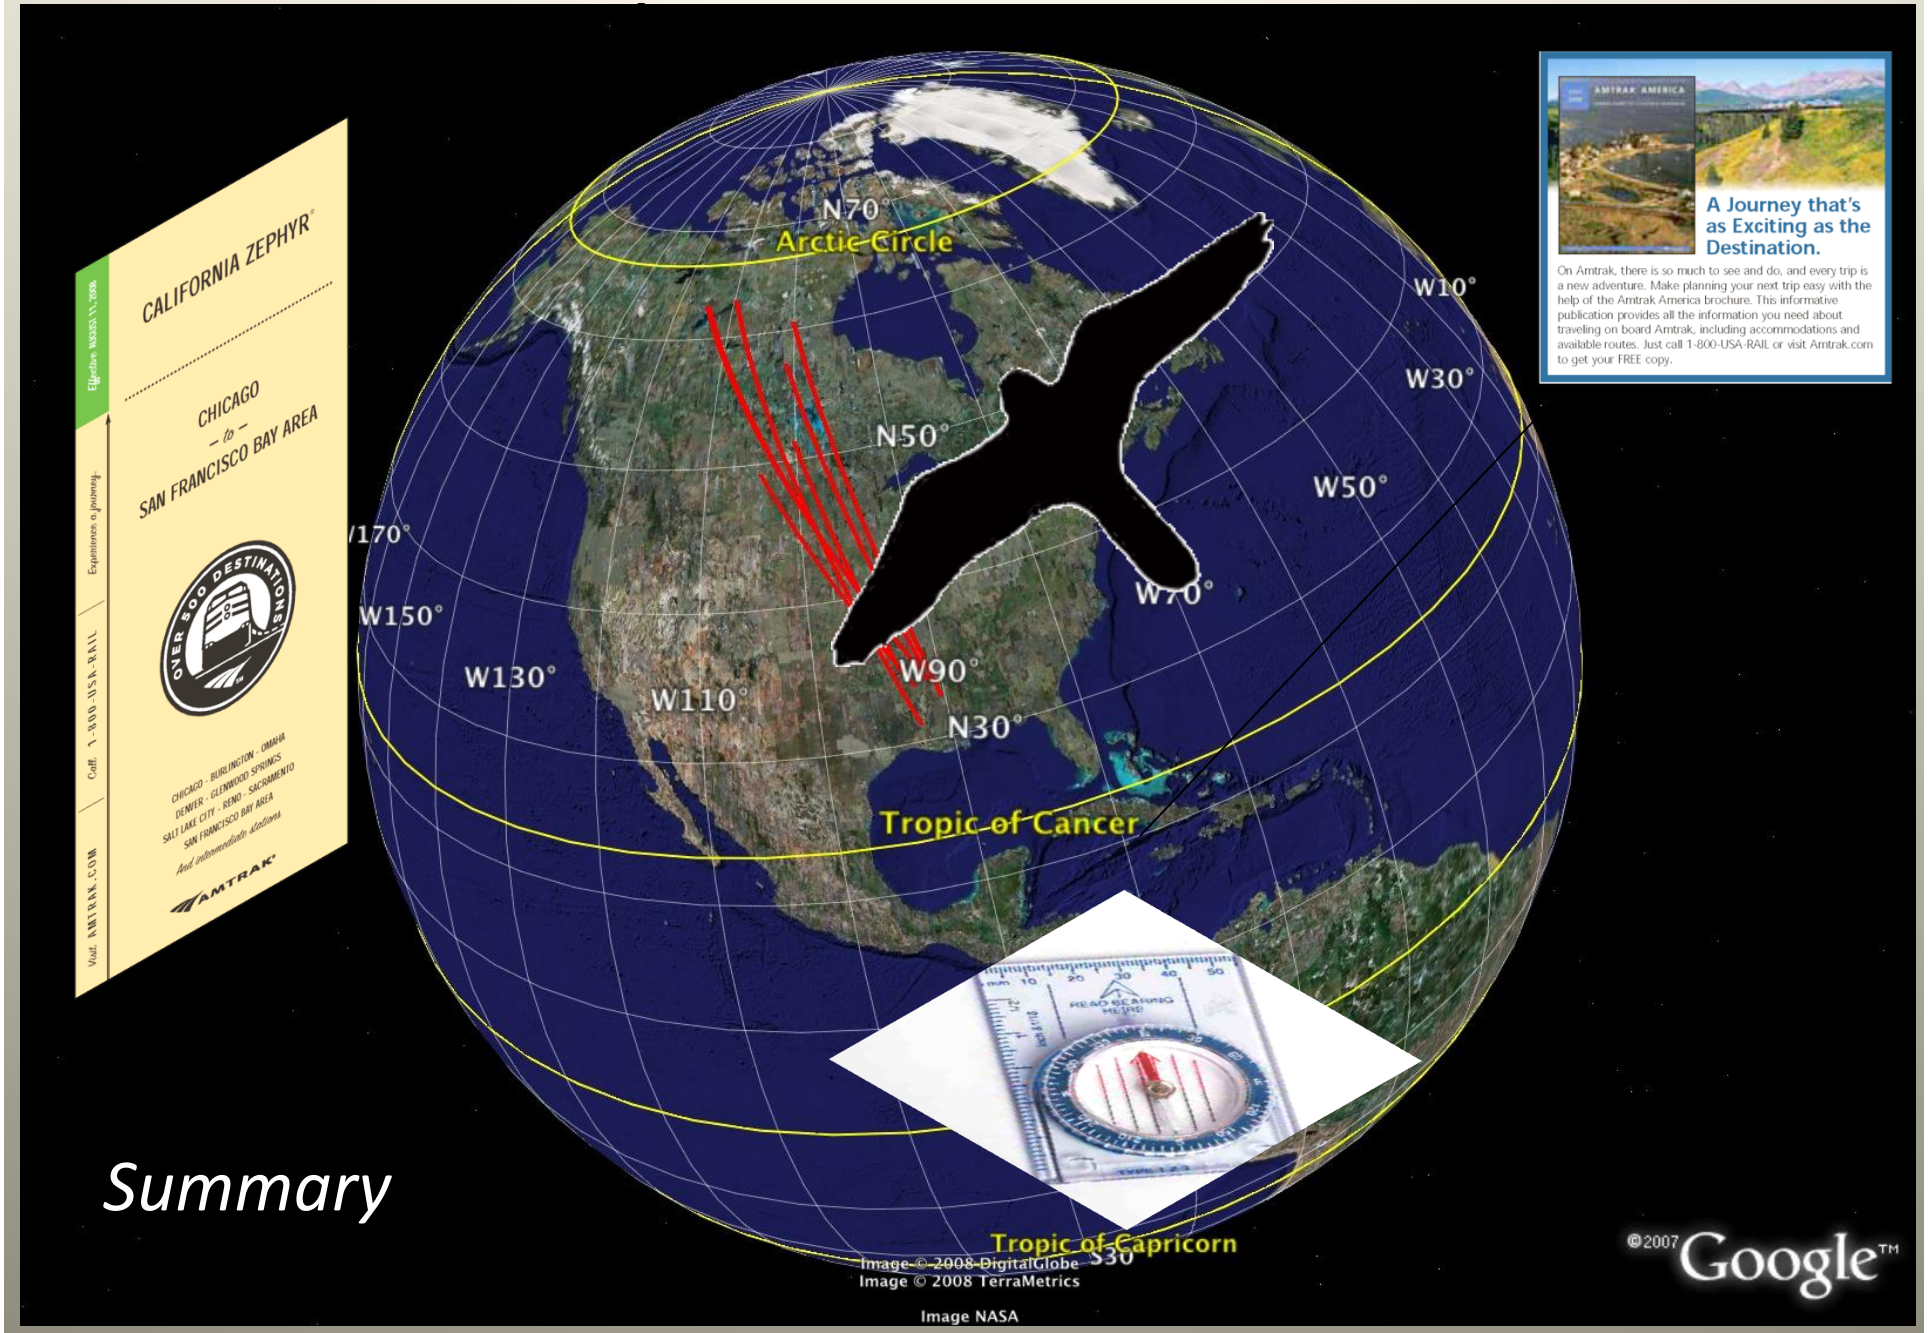
CALIFORNIA ZEPHYR®
CHICAGO
– to –
SAN FRANCISCO BAY AREA
OVER 500 DESTINATIONS
CHICAGO - BURLINGTON - OMAHA
DENVER - GLENWOOD SPRINGS
SALT LAKE CITY - RENO - SACRAMENTO
SAN FRANCISCO BAY AREA
And intermediate stations
AMTRAK
Visit AMTRAK.COM
Call 1-800-USA-RAIL
Experience a journey.
Effective NOVEMBER 19, 2006
N70°
Arctic Circle
N50°
N30°
Tropic of Cancer
Tropic of Capricorn
S30°
W10°
W30°
W50°
W70°
W90°
W110°
W130°
W150°
W170°
A Journey that's as Exciting as the Destination.
On Amtrak, there is so much to see and do, and every trip is a new adventure. Make planning your next trip easy with the help of the Amtrak America brochure. This informative publication provides all the information you need about traveling on board Amtrak, including accommodations and available routes. Just call 1-800-USA-RAIL or visit Amtrak.com to get your FREE copy.
Image © 2008 DigitalGlobe
Image © 2008 TerraMetrics
Image NASA
©2007 Google™

Summary

Software Entropy Catastrophe

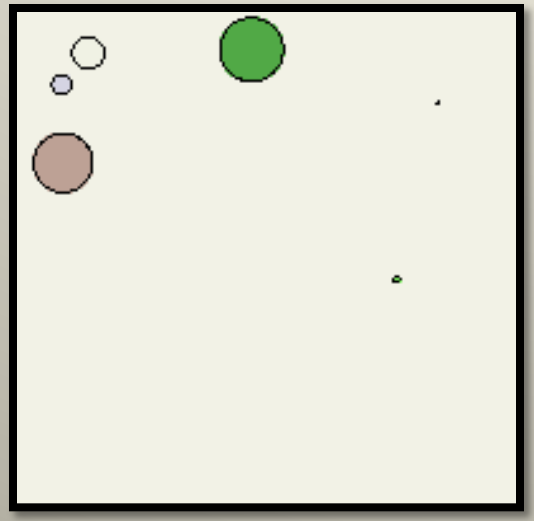

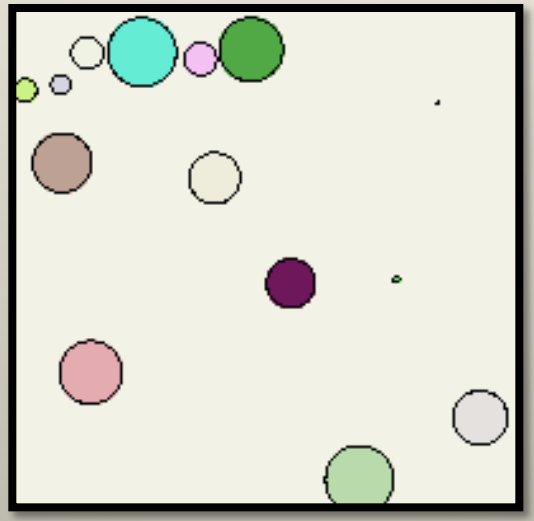

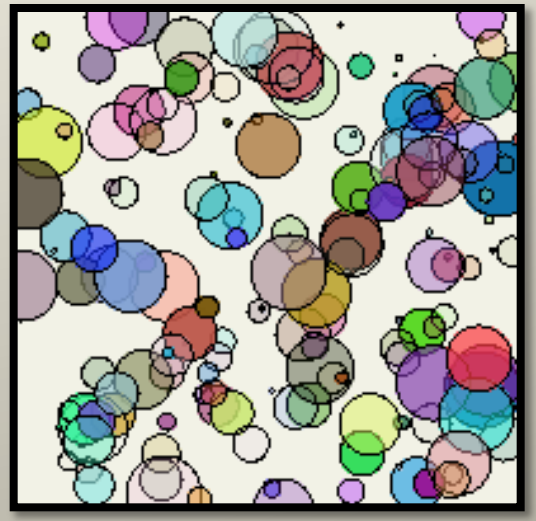

Refactoring ….. Fancy word for

“ Back to Square One “

Individual Based Model

Class “Birds”

Class: Bird

- State:
  - Position
  - Velocity
  - Fuel
- Phenotype:
  - Orientation
  - Stopover

- Update Energy:
  - Resting
  - Feeding
  - Flying
- Update state:
  - Position
  - Velocity
  - Fuel
- Update Time:
  - Resting
  - Feeding
  - Flying

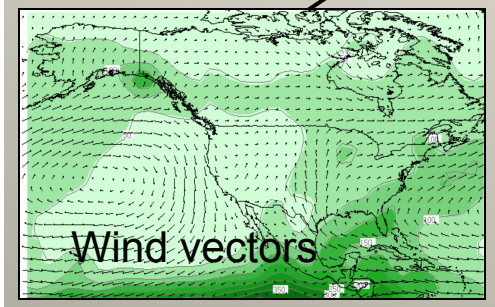

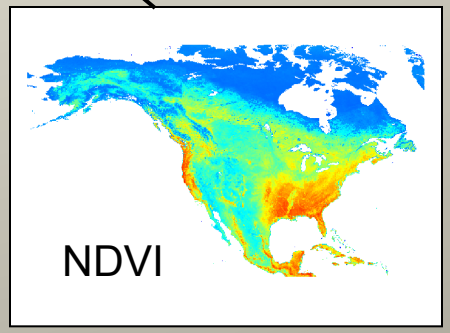

{ Species: Pectoral Sandpiper
Extends Bird }

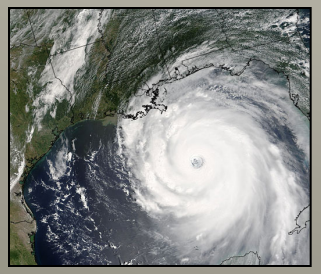

Simulate the migration routes, timing, and energy budgets of individual birds under dynamic weather and land surface conditions

{ Iterate over: Population, daily time step }

Interacting Species

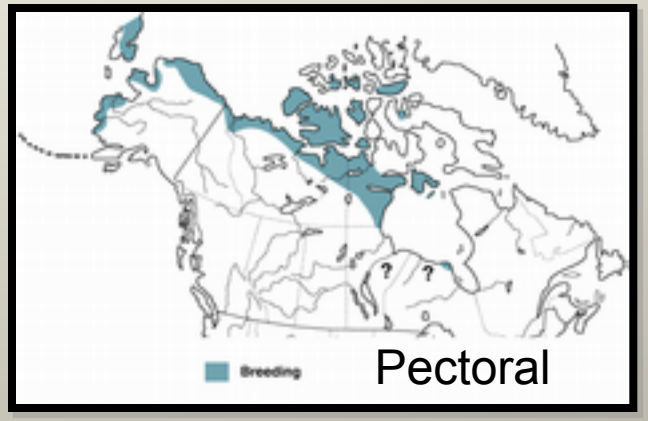


USFWS

Refactoring

Abstract Class: Bird

- State:
- Behavior:
- Update Energy:
- Update Activity:
- Update State:

{ Species: Pectoral Sandpiper Extends Bird }

- Move
- Energy Use

{ Species: Peregrin Falcon Extends Bird }

- Move
- Energy Use

Simulate the migration routes, timing, and energy budgets of individual birds under dynamic weather and land surface conditions

Hardware Architecture

24 GFLOPS

128 Floating Point

NVIDIA

$ 241.43

367 GFLOPS

Graphical Processing Unit

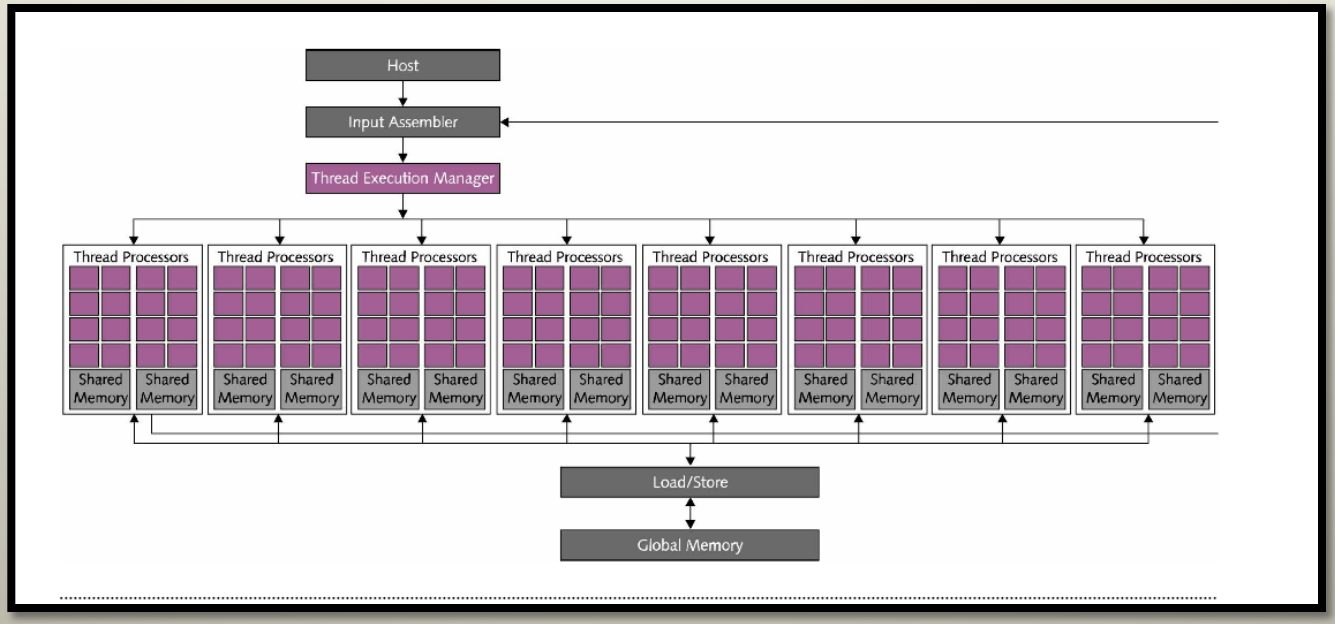

128 floating point thread processing units
(12,228 concurrent threads)

367 GFLOPS , high parallel data transfer

Expect minimum 10X, maybe 30X speedup

Tilmes --- You're not thinking big enuf

Penguin-Tesla 8 TFLOP Workgroup GPU Compute Cluster

Penguin's Tesla cluster bundle provides 8 TFLOPS of parallel GPU computing power in an integrated, racked, ready-to-run package. Penguin's cluster integration expertise and experience with NVIDIA Tesla GPU solutions ensures instant productivity from the time your cluster rolls out from the crate. Our bundled solution also includes Scyld ClusterWare cluster management and provisioning tool which automates both hardware management and job scheduling for distributed environments.

Cluster includes:

- **Altus 1701 Master Node**
  Dual Opteron 2.3GHz Quad-core, 8GB RAM, 1TB drive
- **Altus 1702 Compute Nodes**
  2 "twin" chassis = total 4 nodes! Each with dual Opteron 2.3GHz CPUs and 8GB RAM
- **8 NVidia Tesla GPUs**
  2x Tesla S1070s, each with 4 GPUs. Connected via PCI-E to compute nodes
- **Scyld ClusterWare 5**
  Scyld ClusterWare 5 cluster management software suite. 1 year subscription.
- **InfiniBand DDR Switch**
- **24-port Gigabit Ethernet Switch**
- **15" LCD Display/Keyboard/Mouse Drawer**
- **Rack, PDU, cabling**
- **Rapid delivery**
- **3-year, top-to-bottom warranty on all components. Onsite service also available.**

Resources

Altus 1702 Datasheet
Scyld ClusterWare Datasheet
What is GPU Computing?

$37,995

Penguin Computing has built its strong reputation within the HPC community for superb clustering solutions that enjoy a level of support rarely equaled in our industry.

Vision

Before “refactoring” -- running 10,000 birds on North American Grid

Estimate 400,000 pectoral sandpipers migrating north each season --- likely high mortality (50%?)

(“NASA is an R & D Agency”)

After “refactoring” – creating bird data structures with
70,000,000 (32 bit Java) --- 250,000,000 (64 bit Java)

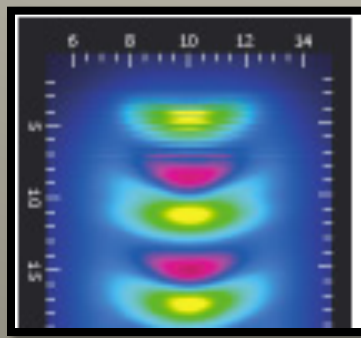

Tech-X Corporation

Tech-X Corporation uniquely combines object-oriented software, distributed technologies, simulation and modeling, and massively parallel computing expertise to assist customers in solving the most difficult scientific problems.

Esias, Miguel, JoAnn, Molly, Ed …

“What are you? --- near sighted ….

“You’re vision is way too narrow ...”

Innovation/IT/Environmental/ Technological/Educational Test Bed

- Compute Platforms
- In-situ wireless networks
- Tasking
- Up-scaling
- Climate change at centers – “System of NASA Centers – “ Prototype at Goddard
- Near by field test bed to work out ideas, students, …
- Wayne --- don’t forget the “birds and the bees”

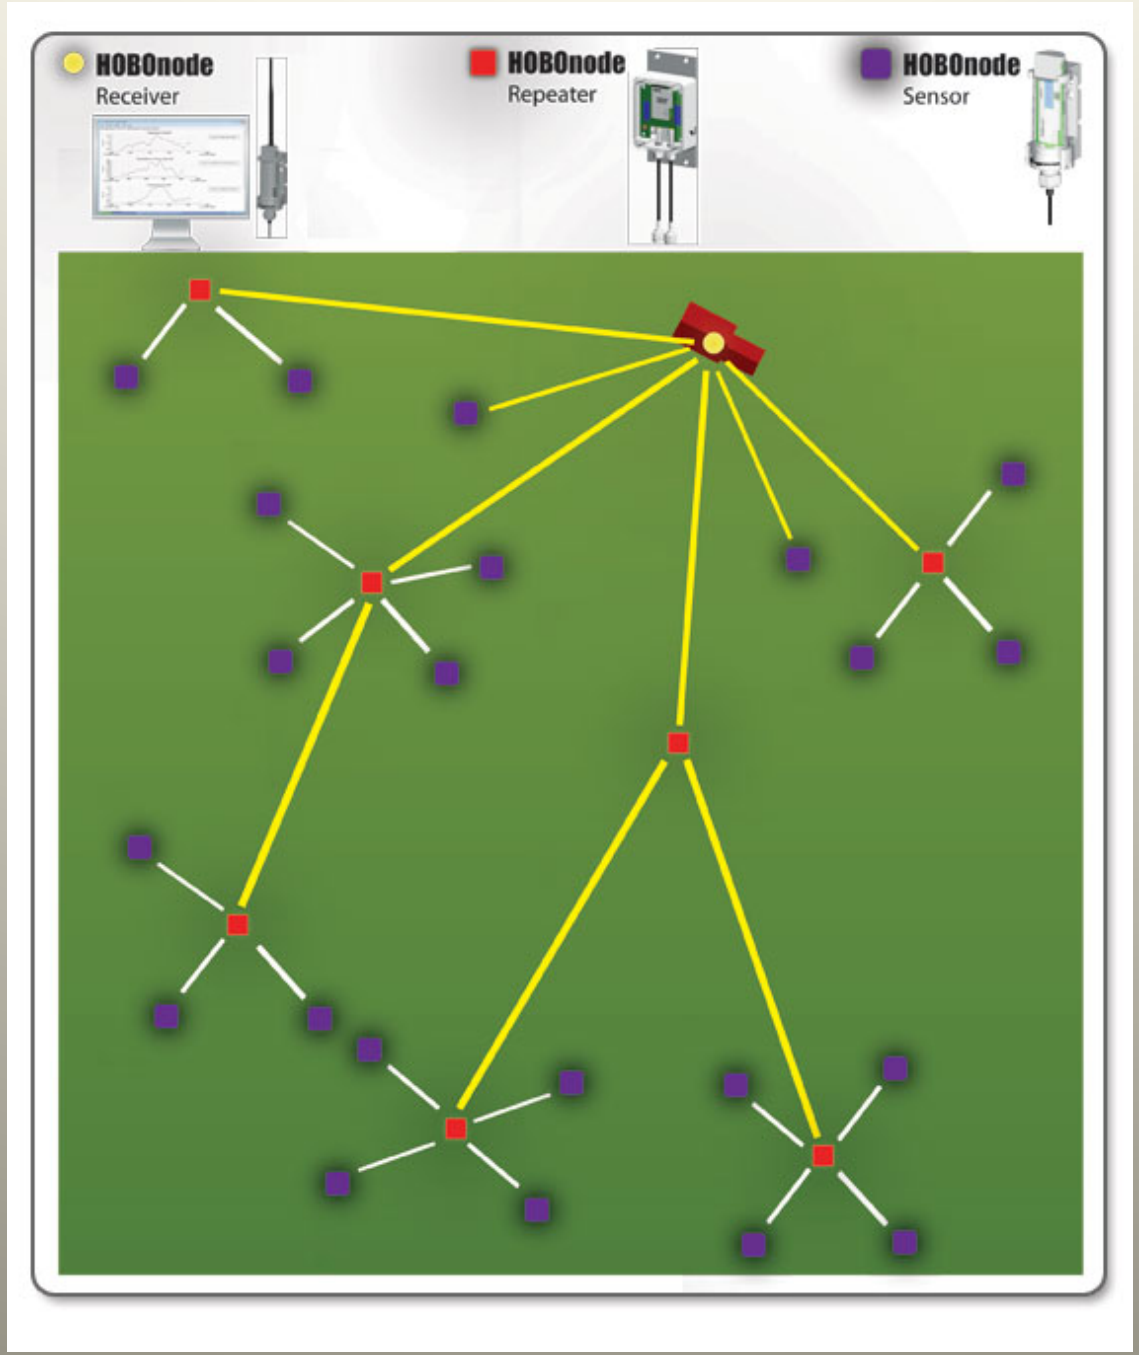
HOBOnode
Receiver
HOBOnode
Repeater
HOBOnode
Sensor

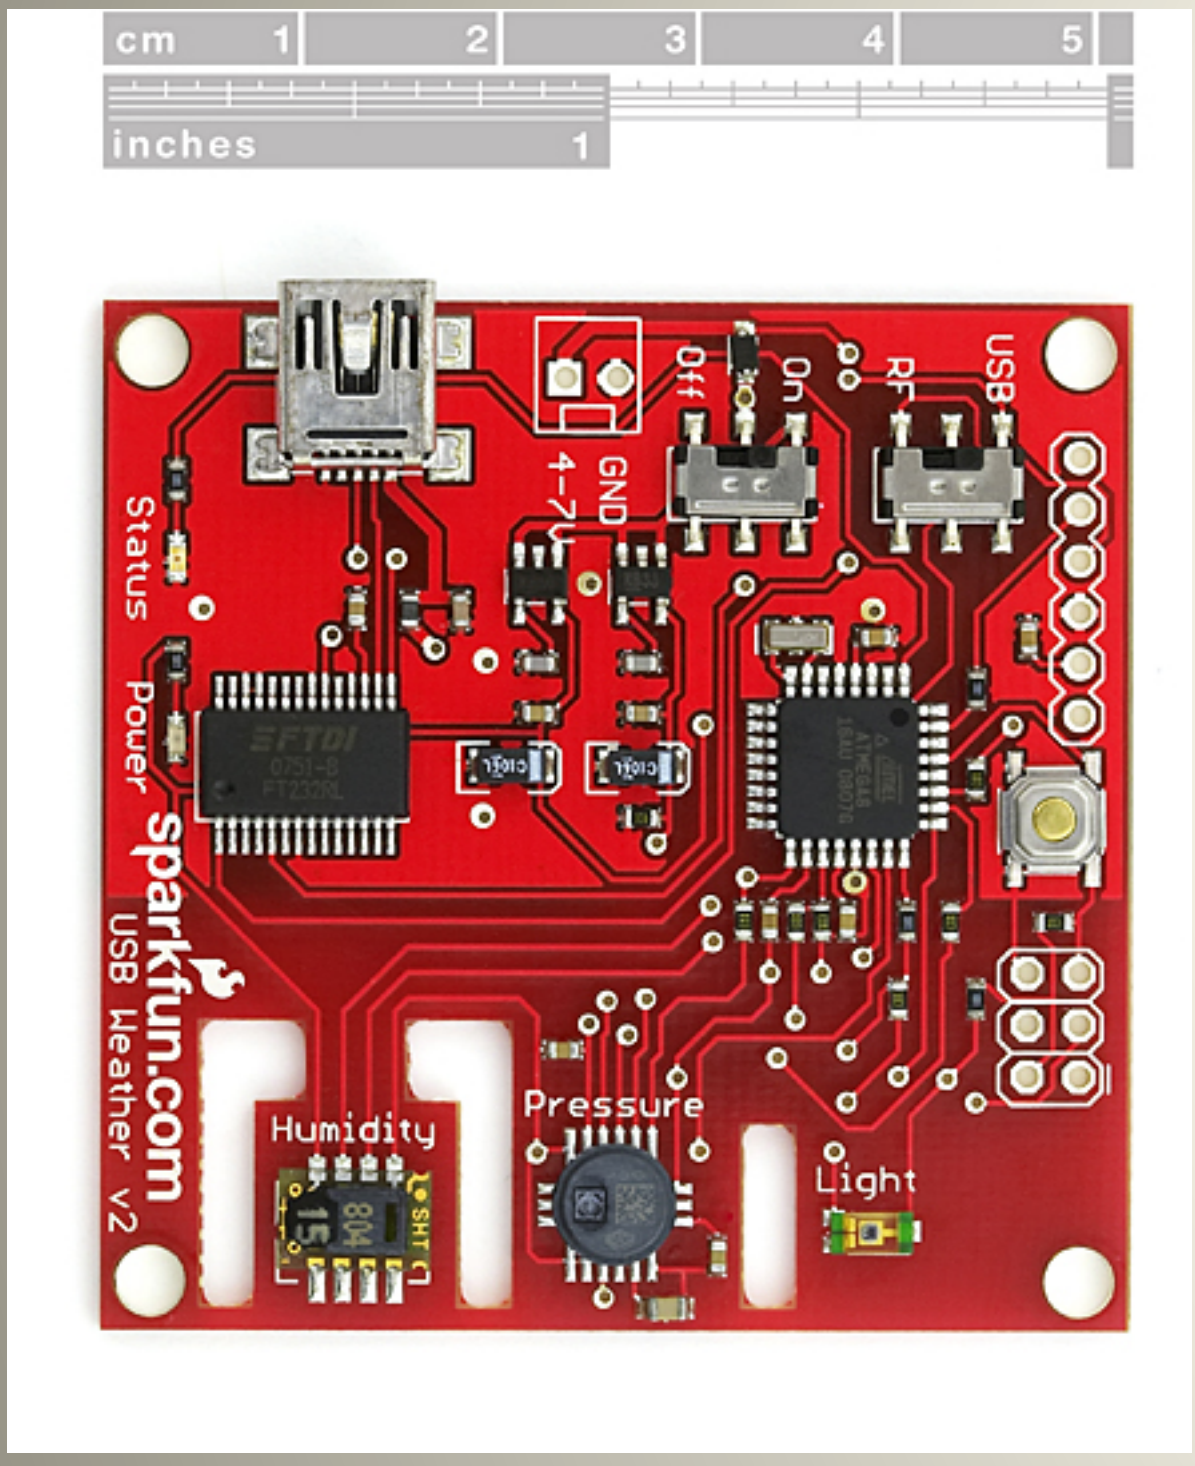
cm 1 2 3 4 5
inches 1
Off
On
RF
USB
GND
4-7V
Status
Power
sparkfun.com
USB Weather v2
Humidity
Pressure
Light
FTDI
FT232RL
SHT

FLIR™
The Global Leader in Infrared Cameras

NEW! FLIR SC660

R&D INFRARED CAMERA SYSTEM

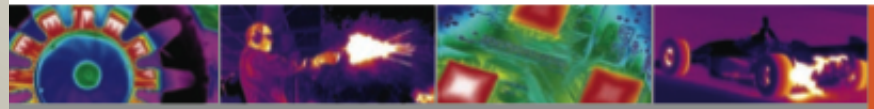

Highest sensitivity and most advanced feature set available. Supplies a combination of infrared and visible spectrum images of superior quality and temperature measurement accuracy – plus GPS, voice annotation, and a host of other advanced features.

- Uncooled 640×480 IR Detector Array
- Thermal Sensitivity ≤45mK
- Built-in 3.2 Mpixel visual camera
- Temperature Range: –40°C to 1500°C
- Full Radiometric Real-time Video to PC
- Automatic GPS Data
- Text and Voice Annotation
- Optional Wireless Remote Operation

Features both thermal and visual camera capabilities – at the touch of a button!

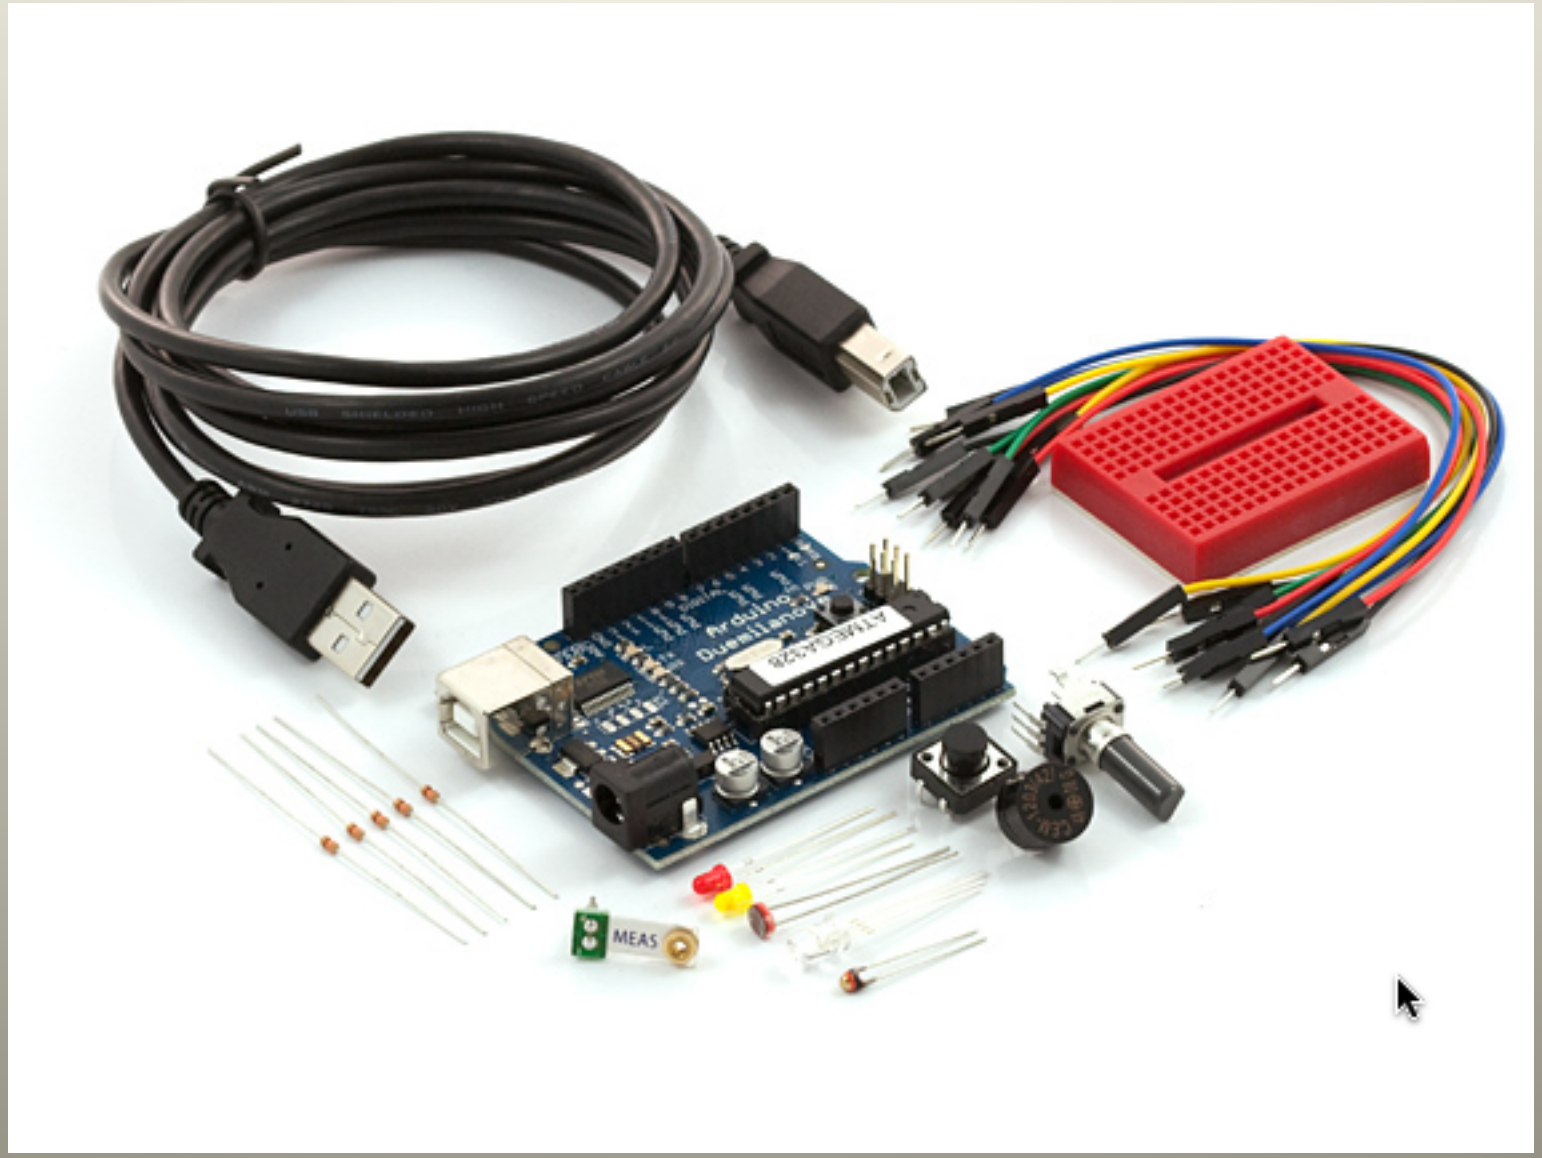
ATMEGA328
MEAS

Wayne Esias

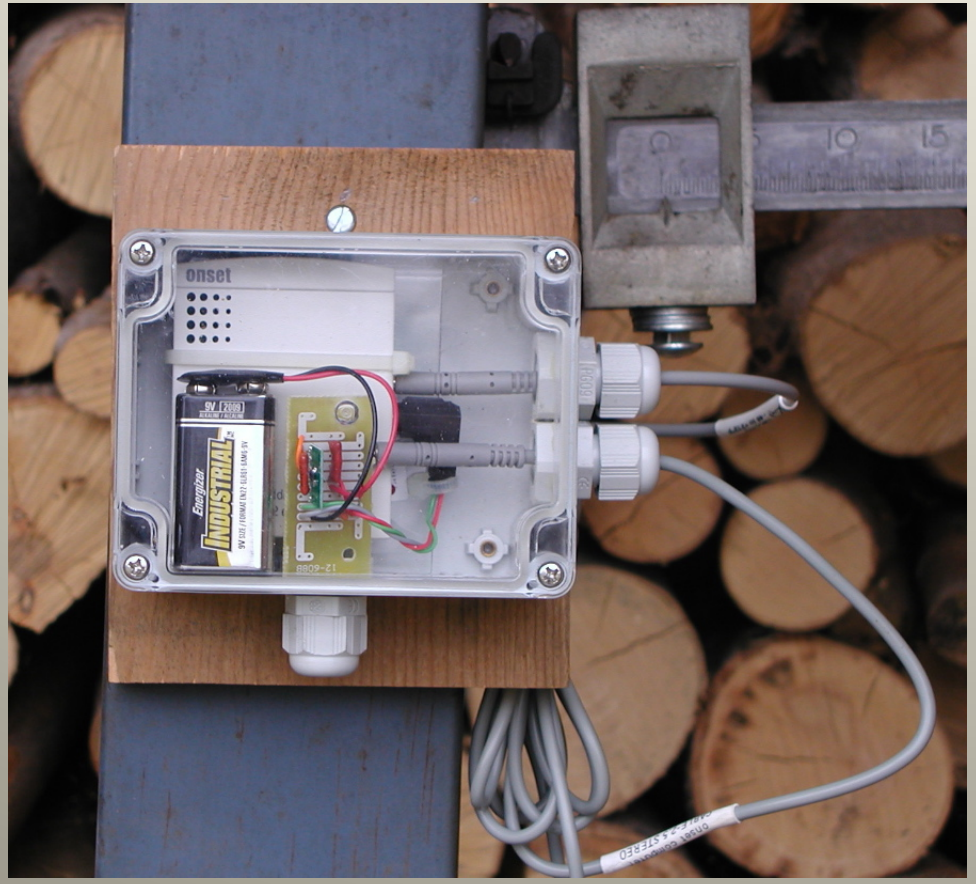

Curt .. .yet again …

Penguin Computing Launches HPC Cloud Computing with GPUs

August 17th, 2009

Penguin Computing has launched a new service that enables high-performance computing within a cloud computing infrastructure, including support for GPU computing with NVIDIA Tesla GPUs. From HPCWire:

> SAN FRANCISCO, Aug. 11 — Penguin Computing, experts in high performance computing solutions, today announced the immediate availability of "Penguin on Demand" — or POD — a new service that delivers, for the first time, a complete high performance computing (HPC) solution in the cloud. POD extends the concept of cloud computing by making optimized compute resources designed specifically for HPC available on demand. POD is targeted at researchers, scientists and engineers who require surge capacity for time-critical analyses or organizations that need HPC capabilities without the expense and effort required to acquire HPC clusters.
>
> POD provides a computing infrastructure of highly optimized Linux clusters with specialized hardware interconnects and software configurations tuned specifically for HPC. Rather than utilizing machine virtualization, as is typical in traditional cloud computing, POD allows users to access a server's full resources at one time for maximum performance and I/O for massive HPC workloads.
>
> Comprising high-density Xeon-based compute nodes coupled with high-speed storage, POD provides a persistent compute environment that runs on a head node and executes directly on the compute nodes' physical cores. Both GigE and DDR high-performance Infiniband network fabrics are available. POD customers also get access to state-of-the-art GPU supercomputing with NVIDIA Tesla processor technology. Jobs typically run over a localized network topology to maximize inter-process communication, to maximize bandwidth and minimize latency.

Ideas

“ Ideas are like the stars,

We never reach them,

but like the mariners

on the sea, we chart

our course by them

Robert H. Smith